On Adaptive Mesh Refinement for Atmospheric Pollution Models

Emil M. Constantinescu and Adrian Sandu

Department of Computer Science, Virginia Polytechnic Institute
and State University, Blacksburg, VA 24061
{emconsta, asandu}@cs.vt.edu

Abstract. This paper discusses an implementation of an adaptive resolution system for modeling regional air pollution based on the chemical transport model STEM. The grid adaptivity is implemented using the generic tool Paramesh. The computational algorithm uses a decomposition of the domain, with the solution in different sub-domains computed at different spatial resolutions. We analyze the parallel computational performance versus the accuracy of long time simulations.

Keywords: Air Pollution Modeling, Adaptive Mesh Refinement.

1 Introduction

Inadequate grid resolution can be an important source of errors in air pollution modeling (APM) where large spatial gradients of tracer concentrations result from the complex interactions between emissions, meteorological conditions, and nonlinear atmospheric chemistry [9].

Chock et. al. [3] studied the effects of grid resolution on model predictions of non-homogeneous atmospheric chemistry. They concluded that increasing the grid size leads to a reduction of the suppression of ozone (O_3) in the presence of high nitrogen oxides ($NO_X = NO + NO_2$), and a decrease in the strength of the NO_X inhibition effect. O_3 loses nearly all the detail near the emission source in the coarse grid case.

A popular multi-resolution approach in air pollution and meteorological modeling is static nesting of finer grids into coarser grids. This approach requires a priori knowledge of where to place the high resolution grids inside the modeling domain; but it does not adjust to dynamic changes in the solution during simulation. In many practical situations the modeler “knows” where higher resolution is needed, e.g. above industrial areas.

In this paper we investigate the parallel performance and accuracy improvements for an application of adaptive mesh refinement (AMR) for modeling regional air pollution. The grid adapts dynamically during the simulation, with the purpose of controlling the numerical spatial discretization error. Unlike uniform refinement, adaptive refinement is more economical. Unlike static grid nesting,

with AMR it is not necessary to specify in advance which areas need higher resolution; what is required is to define a refinement criterion, which is then used by the code to automatically adjust the grid.

The use of generic AMR tools (Paramesh [6, 7]) allows us to harness the power of parallel computers for regional air pollution simulations. Parallel computing is needed as the higher the resolution, the more expensive the simulation becomes.

The paper is organized as follows. Sec. 2 gives an overview of previous work. A brief description of the static mesh APM application and of the AMR system used in this paper is given in Sec. 3. The refinement criterion used in this paper is described in detail in Sec. 3.1. Numerical results are shown in Sec. 4, and Sec. 5 presents conclusions and future research directions.

2 Previous Work

Adaptive meshes have been used in the study of pollutant dispersion in atmosphere [10, 9, 5]. In this section we will discuss some atmospheric AMR applications.

The Ph.D. dissertation of van Loon [10] is focused on numerical methods for smog prediction. The model developed, C_WIROS, has 4 vertical layers and its horizontal domain covers all Europe. The horizontal resolution is 60×60 Km, with 4 levels of refinement. The spatial error estimator uses the curvature of the concentration fields. Specifically, for each column the horizontal curvature is estimated for each species and normalized. The column is flagged for refinement if the error estimator is larger than a user-prescribed tolerance.

Srivastava, McRae and Odman [9] discuss a very interesting approach to grid adaptivity (DSAGA-PPM) for simulating reactive atmospheric pollutants. DSAGA-PPM uses horizontal (2D) adaptivity and employs a constant number of grid nodes. This keeps the total computational time for a simulation manageable. A weight function is defined by a linear combination of curvatures of different chemical species. Based on this weight function the grid is adapted.

3 Implementation Considerations

In this section we describe the numerical and software components of our application. First, we discuss the science application and next we briefly present the AMR approach.

The core science application used in this paper is the state-of-the-art regional APM, STEM [2]. The original code uses a fixed grid, with all the data structures being multidimensional arrays.

STEM solves the advection-diffusion-reaction equation for N species on a 3-D domain:

$$\frac{\partial}{\partial t} c_i + \nabla \cdot (uc_i) = \nabla \cdot (D\nabla c_i) + f(c_i), \text{ with } i = 1 \dots N.$$

The equation is solved using an operator splitting approach. STEM uses linear finite difference discretizations of the transport terms and Rosenbrock methods for solving stiff chemistry [8].

Paramesh offers the infrastructure for AMR on a 2D structured grid. Paramesh is a parallel AMR Fortran toolkit developed by P. MacNeice and K. Olson at the NASA Goddard Space Flight Center [6, 7]. The adaptive resolution is based on a Schwarz-type domain decomposition, with a single Schwarz iteration. We use a two-dimensional (horizontal) grid refinement approach. All the data associated with a column in the original STEM (referred to as STEM variables) are assigned to a mesh-point (cell) in Paramesh, including geographical and meteorological data and species concentrations along the z axis.

The domain is divided into blocks, each containing 6×6 cells plus two guard-cells along each block boundary [1]. At the coarse level, each cell has a resolution of 80×80 Km. For the TraceP (described in Sec. 4.1) simulation over East Asia, the computational domain is covered by 15×10 blocks. At the finest level (level 4) each cell has a resolution of 10×10 Km. Data are linearly interpolated between refinement levels during each mesh refinement/derefinement operation.

During the simulation data available in STEM-specific data types need to be copied into Paramesh data types. The initial species concentrations and geographical information are provided at coarse level at the beginning of the simulation. Meteorological fields, boundary conditions and surface emissions are updated every hour. All data are given at the coarse level, except for the emission inventories which are provided at a fine resolution (10×10 Km).

Experimentally, we noticed a loss in accuracy associated with block refinement near the physical boundary due to the fact that boundary data are available at coarse level only. A refinement restriction was applied to blocks neighboring the domain boundary, such that they are maintained at coarse levels.

The regridding process is handled by Paramesh. Each time the regridding is performed, blocks on each processor are refined or derefined according to one of our criteria and then migrated to other processors if necessary with the goals of load balancing and data locality.

3.1 Refinement Criteria

The estimation of the spatial error in a $NXB \times NYB = N^2$ cells horizontal block (at vertical level k) is done based on the horizontal curvature of the concentration field c at each point (i, j, k) in space

$$err_{i,j,k} = |c_{i+1,j,k} - 2c_{i,j,k} + c_{i-1,j,k}| + |c_{i,j+1,k} - 2c_{i,j,k} + c_{i,j-1,k}|,$$

and by taking the root mean square value normalized by the maximum concentration inside the block

$$ERR_k(c) = \begin{cases} \frac{\sqrt{\sum_{i,j} err_{i,j,k}^2}}{N \cdot \max_{i,j} c_{i,j,k}} & \text{if } \max_{i,j} c_{i,j,k} \geq Atol \\ 0 & \text{if } \max_{i,j} c_{i,j,k} < Atol \end{cases}$$

Note that the error is ignored if the concentration inside the block is below a user-prescribed level, $Atol$. The total error estimate in a column is taken to be the maximum error among all layers $ERR(c) = \max_k ERR_k(c)$. The block is flagged for refinement if $ERR(c) \geq uptol$ and for derefinement if $ERR(c) \leq lowtol$.

The model calculates the concentrations of a large number of trace species and the refinement pattern depends on which concentrations are used for error estimation. We consider a multiple species criterion, focusing on O_3 , formaldehyde ($HCHO$) and NO_X compounds – the main precursors of O_3 . A weighted combination of the regarded species is considered: $w_1 NO + w_2 NO_2 + w_3 O_3 + w_4 HCHO$, with $w_{1,2} = 35\%$ and $w_{3,4} = 15\%$. The error for a mesh-point, based on ℓ species $i_1, \dots, i_\ell$ (in our case $\ell = 4$) the error is estimated by

$$ERR(c_{i_1} \dots c_{i_\ell}) = \sqrt{\frac{1}{\ell} \sum_{j=1}^{\ell} w_j ERR(c_{i_j})^2} \quad .$$

Figure 1.d shows the refined grid pattern corresponding to 0 GMT March 1st, 2001 over East Asia, TraceP conditions with $uptol = 0.25$, $lowtol = 0.1$. The grid is refined in the areas of high emissions, i.e. above industrial regions in China, Japan and Korea. The refinement criteria is applied at simulated hourly increments. In our experiments, we regrid every three hours.

4 Results

In this section we analyze the performance of the parallel implementation for work load and accuracy. Sec. 4.1, describes the experimental setting and Sec. 4.2 discusses the results.

4.1 Experimental Setting

The test problem is a real-life simulation of air pollution in East Asia in support of the TraceP [4] field experiment. The TraceP (NASA TRANsport and Chemical Evolution over the Pacific) field experiment was conducted in East Asia. The meteorological fields, boundary values and emission rates correspond to TraceP starting at 0 GMT of March 4th, 2001 for one week of simulation time. Due to the fact that our initial data are not smooth enough, an imminent transient phase tends to occur if one should start refining from the coarse level. Instead, we simulated two days, starting with March 1st, 2001, at the finest level and after that we applied the mesh refinement criterion and allowed the grid to coarsen. The accuracy and timing results were taken into consideration after three days of simulation (two at the finest level and one for system adjusting to a relative steady state).

The simulated region covers 7200×4800 Km. At the coarse level (each cell with 80×80 Km) there are 150 blocks, each containing 6×6 cells. At the finest level, level 4 (10×10 Km) there are 5058 working blocks (182,088 mesh points). Each cell holds a column of 18 layers with 2340 of STEM variables.

The simulations are performed on Virginia Tech’s System X, the fastest academic supercomputer in the world. It has 1100 node Apple XServe G5 dual processor, 4 Gb of RAM per node. The interconnect consists of InfiniBand switches (primary) and Cisco 4506 Gigabit Ethernet (secondary). We were unable to accurately measure execution times for long simulations but we were able to make an estimation of the scalability based on short one simulated hour runs.

4.2 Numerical/Parallel Performance

The timing results for 1 simulated hour at the finest refinement level on 16, 32, 64 and 96 processors are presented in table 1.a. Considering the fact that the workload remains constant, the speed-up is relatively good especially when using 1 processor per node. The computational intensive part of our application shows an insignificant improvement when we switch from one to two processors per node. On the other hand, the communication intensive portion shows a large improvement when switching from two to one processors per node. The reason for that is probably a less congested communication pattern.

Table 1.b shows the wall-clock for several scenarios (Fine, Coarse and two AMR runs) for one week of simulation. The application tuning for specific processor workload is a problem in itself, especially for parallel implementations, due to the difficulty in managing the amount of refinement that each processor does. Scenario AMR 1 is close to a quarter of the total fine wall-clock and close to our expectations in terms of accuracy as it will be shown below. AMR-2 is very competitive in terms of timing but the accuracy of the simulation is degraded (see Figs. 1.a and 2.a). In our experiments we noticed that accuracy is tightly linked to the number of mesh-points that are concentrated in the higher estimated truncation error locations.

The accuracy results are computed as the error mean for all 18 layers represented as error level contours. The error levels for O_3 after one week of simulation are shown in Figure 1.{a,b,c} for the two AMR results compared to the coarse simulation. The same results are also shown for the NO species in Figure 2.{a,b,c}. AMR-1 has a very high accuracy performance for both species, while AMR-2 has not performed so well. This suggests that insufficient refinement does not bring any significant gains in terms of accuracy. The mesh-point

Table 1. (a) - The wall-clock for one hour of simulated time on the finest refinement level when using one or two processors per node; (b) - Timing for fine, coarse and two AMR cases for one simulated week

No. of Procs.	Time [s]	Time [s]	Simulation type	Time [s]	Final (Mean) no. of mesh-points
	1 proc./node	2 proc./node			
16	2163	2739	Fine	429,299	$36 \times 5058(5058)$
32	1125	1270	AMR-1	126,697	$36 \times 2310(2548)$
64	841	1206	AMR-2	20,526	$36 \times 375(532)$
96	502	816	Coarse	4,150	$36 \times 150(150)$

(a)

(b)

Fig. 1. Ozone error contours (percent) after one week of simulation ending at 0 GMT of March 11th, 2001, for: (a) – AMR-2, (b) – AMR-1 and (c) – coarse level simulation. (d) – The refined grids at (0 GMT of March 1st, 2001) for East Asia during the TraceP campaign. Each block (shown) consists of 6×6 computational cells (not shown). The criterion is based on the curvature of NO_X using $uptol = 0.25$ and $lowtol = 0.1$ with maximum refinement level 4 (10×10 Km)

dynamics over the one week period is shown in Figure 2.d. As we have expected, the system finds itself in a relative steady state - fine grids may move but the overall number of mesh-points is kept relatively constant, decreasing slowly as the solution becomes more and more smooth.

5 Conclusions

In this paper we investigate the parallel performance and accuracy improvements of an adaptive grid air pollution model based on the parallelized STEM air pollution model with the Paramesh tool. We look for accuracy improvements in O_3 and NO_X , species and low computational overheads.

The model scales very well as long as we use one processor per node. The communication intensive part plays a very important role as we tested our application

Fig. 2. NO error contours (percent) after one week of simulation ending at 0 GMT of March 11th, 2001, for: (a) – AMR-2, (b) – AMR-1 and (c) – coarse level simulation. (d) – Represents the evolution of the number of blocks over the simulated week

on a network of workstations. This may be alleviated by a better mesh-point to processor partitioning.

The workload corresponding to the number of mesh-points is determined by the refinement criterion's refinement and derefinement thresholds. We found it very difficult to find appropriate values for those tolerances due to the relative autonomy of each process. A possible mitigation for this problem would be to collect all the mesh-point truncation errors, rank them and allow refinement/derefinement to a limited number of blocks. We chose to let the system evolve without any restrictions.

The accuracy benefits of AMR are amplified by the use of large number mesh-points: the finer the mesh, the better the accuracy. In our experiments, the use of almost triple the number of mesh-points than for the coarse simulation did not bring significant accuracy improvements to the final solution, while keeping the number of mesh-points between one quarter and one third of the fine one showed just a little solution degradation.

The dominant errors are located downwind of the emission sources especially for the high resolution AMR simulations (large number of mesh-points). A possible explanation is that the effect of errors in regions with high emissions are amplified by the chemical processes and advected downwind. This aspect would suggest a refined grid (increased resolution) upwind of the area of interest.

Acknowledgements

This work was supported by the National Science Foundation through the awards NSF CAREER ACI 0093139 and NSF ITR AP&IM 0205198. Our special thanks go to Virginia Tech's TCF for the use of the System X cluster.

References

1. C. Belwal, A. Sandu, and E. Constantinescu. Adaptive resolution modeling of regional air quality. *ACM Symposium on Applied Computing*, 1:235–239, 2004.
2. G.R. Carmichael. STEM – A second generation atmospheric chemical and transport model. URL: <http://www.cgrer.uiowa.edu>, 2003.
3. D.P. Chock, S.L. Winkler, and P. Sun. Effect of grid resolution and subgrid assumptions on the model prediction of non-homogeneous atmospheric chemistry. *The IMA volumes in mathematics and its applications: Atmospheric modeling*, D.P. Chock and G.R. Carmichael editor, pages 81–108, 2002.
4. G.R. Carmichael et. al. Regional-scale chemical transport modeling in support of the analysis of observations obtained during the TRACE-P experiment. *J. Geophys. Res.*, 108:10649–10671, 2004.
5. S. Ghorai, A.S. Tomlin, and M. Berzins. Resolution of pollutant concentrations in the boundary layer using a fully 3D adaptive technique. *Atmospheric Environment*, 34:2851–2863, 2000.
6. P. MacNeice and K. Olson. PARAMESH V2.0 – Parallel Adaptive Mesh Refinement. URL: http://ct.gsfc.nasa.gov/paramesh/Users_manual/amr.html, 2003.
7. P. MacNeice, K. Olson, and C. Mobarry. PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126:330–354, 2000.
8. A. Sandu, Dacian N. Daescu, Gregory R. Carmichael, and Tianfeng Chai. Ad-joint sensitivity analysis of regional air quality models. *Journal of Computational Physics*, :Accepted, 2004.
9. R.K. Srivastava, D.S. McRae, and M.T. Odman. Simulation of a reacting pollutant puff using an adaptive grid algorithm. *Journal of Geophysical Research*, 106(D20):24,245–24,257, 2001.
10. M. van Loon. *Numerical Methods in Smog Prediction*. Ph.D. Dissertation, CWI Amsterdam, 1996.